



CD-doc-666-v5.doc  
GCC Automatic Shutdown Plan

David J. Ritchie  
Operations, Computing Division

February 28, 2005

## Introduction

The purpose of this document is to describe the plan for providing an automatic shutdown capability for the analysis computing in GCC.

## Background

GCC, the Grid Computing Center, sited within the building formerly known as "Wideband," will house a large number of PCs (up to 2800) performing high energy physics computations and connected to the laboratory network through a number (~5) of subnets. The PCs, which generally run Linux, are powered via a 1000 KVA UPS system which obtains *its* power from the laboratory electrical feeders supplied by Commonwealth Edison via the laboratory substation. Cooling is provided through a number (~9) of Liebert CRAC units.

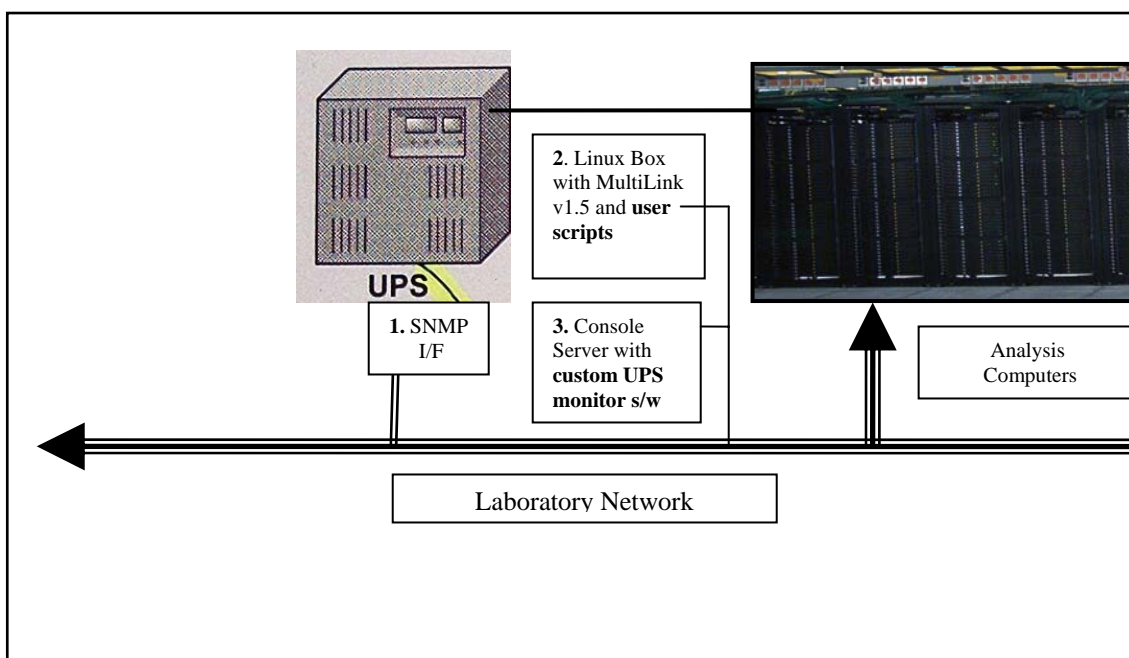
The purpose of the UPS is to make it possible for the PC's to shutdown in a controlled way (i.e., "a soft landing") in the event of a failure of the laboratory power. Experience has shown that providing a soft landing will minimize hardware problems that result from unexpected power outages.

The capacity of the UPS makes it possible for a totally installed GCC (i.e., 2800 PCs) to have power for a maximum of 13 minutes. The intent, of course, is to have the PCs shutdown in two to five minutes.

Note that the CRACs are not powered through the UPS but instead use laboratory power. It is not a goal to keep the CRACs powered during the soft shutdown process – instead, it is intended to shut the PCs down before the temperature rises to a point where it would be dangerous to the equipment.

## Plan for Implementing Automatic Shutdown at GCC

The plan is to purchase the MultiLink capability as diagrammed below:



### Components as numbered above:

1. UPS in GCC with SNMP I/F card (~\$500) and MultiLink Network Shutdown License (~\$250) and MultiLink v1.5 software provided on CDROM media (~\$50). This will allow software to detect UPS state: on-battery, low battery, return to normal (i.e., on utility power), and weak battery.
2. Linux PC running the MultiLink v1.5 Client located in GCC Computing Room with user-written scripts to shutdown other nodes:
  - a. One or more needed for CMS.
  - b. One or more for Run-II.
3. Console Server provided by Farms Group with custom UPS monitoring software provided by Farms Group.

## Cost Summary

Item	Amount
SNMP I/F Card for UPS (est.)	\$500.00
One MultiLink™ 1.5 Shutdown software. (Provided on CDROM).	50.00
One MultiLink™ network shutdown license for up to 10 computers.	\$250.00
<b>Total</b>	<b>\$800.00</b>

## Advantages Summary

1. The UPS is well monitored using manufacturer supported instrumentation.
2. There are two alert points so that implementers of scripts have some options as to whether to trigger on the on-batteries event or on the os-shutdown event. This makes it possible, for example, to interpret the first as directed towards the analysis programs while the second as directed towards the power controllers.

## Extensibility Summary

The extensibility of this automatic shutdown solution allows:

1. Extension of this to additional nodes is possible.

It is noted that the SNMP card is stated to be limited to no more than twenty computers interacting with it. At present, it is anticipated that it is interacting with eight so this seems well within that limit.

2. Extension of the same system to monitor from FCC is also possible.

## Difficulties Overcome

The SNMP protocol is used only for monitoring—not control so that concerns about the vulnerabilities of the v1 protocol are overcome.

The different ways of interacting with the UPS signals desired by the groups responsible for the analysis computers are accommodated.

### **Alternatives Considered**

We have considered a number of alternatives to the above automatic shutdown plan (Thanks to Keith Chadwick for bringing these to our attention):

1. Netbotz—(<http://www.netbotz.com>) We have looked at the Netbotz management and monitoring device option. While this option has transmission of alerts via SSL and so avoids the SNMPv1 problem, it does not have the capability as near as I can tell of executing a script on the Head Node when an alert occurs.
2. Omnitronix—(<http://www.omnitronix.com>) We have examined the information on the Omnitronix web site. Its offerings make use of SNMP extensively. In particular, it does have AlarmManager software but it only apparently on the Windows operating system. We need the ability to send alarms to Linux boxes, however, so this is ruled out.